

RESULT LIST

8 results found in the Worldwide database for:
snapshot in the title AND **metadata** in the title or abstract
(Results are sorted by date of upload in database)

Data supplied from the [esp@cenet](http://esp.cenet.org) database - Worldwide

Inventor: SAWDON WAYNE A; SCHMUCK FRANK B **Applicant:** IBM
EC: **IPC:** G06F11/14
Publication info: AU2003214039 - 2003-09-04

Data supplied from the **esp@cenet** database - Worldwide

SYSTEM AND METHOD FOR PERFORMING AN IMAGE LEVEL SNAPSHOT AND FOR RESTORING PARTIAL VOLUME DATA

Patent number: WO2005048085
Publication date: 2005-05-26
Inventor: VARADHARAJAN PRAKASH (IN); PAWAR RAHUAL (IN); KUMAR AVINASH (IN); PRAHLAD ANAND (US); NGO DAVID (US)
Applicant: COMMVAULT SYSTEMS INC (US); VARADHARAJAN PRAKASH (IN); PAWAR RAHUAL (IN); KUMAR AVINASH (IN); PRAHLAD ANAND (US); NGO DAVID (US)
Classification:
- **International:** (IPC1-7): G06F
- **European:**
Application number: WO2004US38455 20041115
Priority number(s): US20030519576P 20031113; US20030519876P 20031113

Also published as:
 WO2005050386 (A2)

[Report a data error here](#)

Abstract of WO2005048085

The present invention relates to a method for performing an image level copy of an information store. The present invention comprises performing a snapshot of an information store that indexes the contents of the information store, retrieving data associated with the contents of the information store from a file allocation table, copying the contents of the information store to a storage device based on the snapshot, and associating the retrieved data with the copied contents to provide file system information for the copied contents.

Data supplied from the esp@cenet database - Worldwide

RESULT LIST

12 results found in the Worldwide database for:
snapshot in the title AND **address** in the title or abstract
(Results are sorted by date of upload in database)

1 On-line data backup method based on data volume snapshot

Inventor: LI LIHONG (CN); WU JIANG (CN); (+1) Applicant: LENOVO BEIJING CO LTD (CN)

EC: IPC: G06F12/16

Publication info: CN1567262 - 2005-01-19

2 THIN-PROVISIONING WITH SNAPSHOT TECHNOLOGY

Inventor: CAMERON DOUGLAS J (US) Applicant: 3PARDATA INC (US); CAMERON DOUGLAS J (US)

EC: IPC: G06F12/00

Publication info: WO2004102391 - 2004-11-25

3 Fibre channel fabric snapshot service

Inventor: KAUSHIK BALAKUMAR N (US);
BALASUBRAMANIAN SHANKAR (US); (+1) Applicant: BROCADE COMM SYSTEMS INC (US)

EC: H04L29/08A7 IPC: G06F12/00

Publication info: US2004220971 - 2004-11-04

4 Organization of multiple snapshot copies in a data storage system

Inventor: ARMANGAU PHILIPPE (US); BERGANT MILENA Applicant:
(US); (+3)

EC: IPC: G06F17/30

Publication info: US2004030727 - 2004-02-12

5 Writable file system snapshot with ditto address feature

Inventor: SAWDON WAYNE A (US); SCHMUCK FRANK B Applicant: IBM (US)
(US); (+1)

EC: G06F17/30F IPC: G06F7/00

Publication info: US2003158834 - 2003-08-21

6 File system snapshot with ditto address feature

Inventor: HASPIN ROGER L (US); SAWDON WAYNE A Applicant: IBM (US)
(US); (+2)

EC: G06F17/30F IPC: G06F12/00

Publication info: US2003158863 - 2003-08-21

**7 Copy method supplementing outboard data copy with previously
instituted copy-on-write logical snapshot to create duplicate consistent
with source data as of designated time**

Inventor: KACZMARSKI MICHAEL ALLEN (US); WARREN Applicant:
DONALD PAUL (US)

EC: IPC: G06F12/00

Publication info: US2003140070 - 2003-07-24

**8 System and method for remotely creating a physical memory snapshot
over a serial bus**

Inventor: VACHON ANDRE F (US); Applicant:
CHRYSANTHAKOPOULOS GEORGIOS (US)

EC: IPC: G06F11/26

Publication info: US2002078404 - 2002-06-20

**9 Snapshot and recall based mechanism to handle read after read
conflict**

Inventor: GUPTA C NAGESH (US) Applicant: HEWLETT PACKARD CO (US)

EC: G06F13/42C1S IPC: G06F13/14; G06F12/00

Publication info: US6449673 - 2002-09-10

10 System and method for real-time data backup using snapshot copying

with selective compaction of backup data

Inventor: FRANKLIN CHRIS (US)

Applicant: ADAPTEC INC (US)

EC:

IPC: G06F12/00; G06F12/16

Publication info: US6061770 - 2000-05-09

Data supplied from the **esp@cenet** database - Worldwide

RESULT LIST

21 results found in the Worldwide database for:
snapshot in the title AND **disk** in the title or abstract
(Results are sorted by date of upload in database)

1 DISK ARRAY DEVICE WITH SNAPSHOT SIMULATION FUNCTION

Inventor: SAKAI TOSHIHIRO

Applicant: NIPPON ELECTRIC CO

EC:

IPC: G06F3/06

Publication info: JP2005128590 - 2005-05-19

2 BLOCK LEVEL DATA SNAPSHOT SYSTEM AND METHOD

Inventor: GUTHRIE II JOHN L (US)

Applicant: ZETTA SYSTEMS INC (US); GUTHRIE II JOHN L (US)

EC:

IPC: G06F

Publication info: WO2005052734 - 2005-06-09

3 SNAPSHOT QUICKENING METHOD

Inventor: HARUMA YUMIKO; SUGIYAMA HIROHARU;

Applicant: HITACHI LTD

(+3)

EC:

IPC: G06F3/06

Publication info: JP2005055972 - 2005-03-03

4 System and methods for transforming data from a source to target platform using snapshot

Inventor: ANDRE JEFFREY (US); TOMSULA PATRICK J (US)

Applicant: STORAGE TECHNOLOGY CORP (US)

EC:

IPC: G06F15/16

Publication info: US6874035 - 2005-03-29

5 Snapshot marker

Inventor: LAM WAI (US)

Applicant:

EC: G06F12/08B2

IPC: G06F12/08; G06F12/16

Publication info: US2005005070 - 2005-01-06

6 THIN-PROVISIONING WITH SNAPSHOT TECHNOLOGY

Inventor: CAMERON DOUGLAS J (US)

Applicant: 3PARDATA INC (US); CAMERON DOUGLAS J (US)

EC:

IPC: G06F12/00

Publication info: WO2004102391 - 2004-11-25

7 Persistent Snapshot Methods

Inventor: WELSH ALAN L (US); TOLPIN RICHARD M (US); (+8)

Applicant: COLUMBIA DATA PRODUCTS INC (US)

EC:

IPC: G06F12/16

Publication info: US2004117572 - 2004-06-17

8 System and method of an efficient snapshot for shared large storage

Inventor: KIM YOUNG HO (KR); KANG DONG JAE (KR); Applicant: (+4)

EC: G06F11/14A4B1M8

IPC: G06F12/16

Publication info: US2004083345 - 2004-04-29

9 Snapshot creating method and apparatus

Inventor: HARA JUNICHI (JP); NAKANO TAKAHIRO (JP); Applicant: (+3)

EC:

IPC: G06F17/00; G06F7/00

Publication info: US2004103104 - 2004-05-27

10 Persistent Snapshot Management System

Inventor: GREEN ROBBIE A (US); MUIRRAGUI PATRICIO R (US); (+8)

Applicant:

EC:

IPC: G06F12/12

Publication info: US2003167380 - 2003-09-04

Data supplied from the esp@cenet database - Worldwide

RESULT LIST

2 results found in the Worldwide database for:
inode in the title AND **snapshot** in the title or abstract
(Results are sorted by date of upload in database)

1 System and method for asynchronous mirroring of snapshots at a destination using a purgatory directory and inode mapping

Inventor: MANLEY STEPHEN L (GB); OWARA SHANE S Applicant:

(US)

EC: G06F11/20L4M8A; G06F17/30F IPC: G06F12/00

Publication info: **US2003195903** - 2003-10-16

2 System and method for asynchronous mirroring of snapshots at a destination using a purgatory directory and inode mapping

Inventor: MANLEY STEPHEN L (GB); OWARA SHANE S Applicant:

(US)

EC: G06F11/20L4M8A; G06F17/30F IPC: G06F12/00

Publication info: **US2003182325** - 2003-09-25

Data supplied from the **esp@cenet** database - Worldwide

 [Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

file mirroring inode ditto disk address



 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [file mirroring inode ditto disk address](#)

Found 23,140 of 169,166

Sort results by

 [Save results to a Binder](#)

[Try an Advanced Search](#)

Display results

 [Search Tips](#)
 [Open results in a new window](#)

[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale 

1 [Improving storage system availability with D-GRAID](#)

 Muthian Sivathanu, Vijayan Prabhakaran, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

May 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(700.30 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present the design, implementation, and evaluation of D-GRAID, a gracefully degrading and quickly recovering RAID storage array. D-GRAID ensures that most files within the file system remain available even when an unexpectedly high number of faults occur. D-GRAID achieves high availability through aggressive replication of semantically critical data, and fault-isolated placement of logically related data. D-GRAID also recovers from failures quickly, restoring only live file system data to a h ...

Keywords: Block-based storage, Disk array, RAID, fault isolation, file systems, smart disks

2 [The logical disk: a new approach to improving file systems](#)

 Wiebren de Jonge, M. Frans Kaashoek, Wilson C. Hsieh

December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(1.55 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Logical Disk (LD) defines a new interface to disk storage that separates file management and disk management by using logical block numbers and block lists. The LD interface is designed to support multiple file systems and to allow multiple implementations, both of which are important given the increasing use of kernels that support multiple operating system personalities. A log-structured implementation of LD (LLD) demonstrates that LD can be implemented efficiently. LLD adds about 5% to 10% ...

Keywords: MINIX, UNIX, disk storage management, file system organization, file system performance, high write performance, log-structured file system, logical disk

3 Serverless network file systems

 Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe, David A. Patterson, Drew S.

Roselli, Randolph Y. Wang

February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1

Publisher: ACM Press

Full text available:  pdf(2.69 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We propose a new paradigm for network file system design: serverless network file systems. While traditional network file systems rely on a central server machine, a serverless system utilizes workstations cooperating as peers to provide all file system services. Any machine in the system can store, cache, or control any block of data. Our approach uses this location independence, in combination with fast local area networks, to provide better performance and scalability th ...

Keywords: RAID, log cleaning, log structured, log-based striping, logging, redundant data storage, scalable performance

4 Serverless network file systems

 T. E. Anderson, M. D. Dahlin, J. M. Neefe, D. A. Patterson, D. S. Roselli, R. Y. Wang

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95**, Volume 29

Issue 5

Publisher: ACM Press

Full text available:  pdf(2.48 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

5 Doubly distorted mirrors

 Cyril U. Orji, Jon A. Solworth

June 1993 **ACM SIGMOD Record , Proceedings of the 1993 ACM SIGMOD international conference on Management of data SIGMOD '93**, Volume 22 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.05 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Traditional mirrored disk systems provide high reliability by multiplexing disks. Performance is improved with parallel reads and shorter read seeks. However, writes must be performed by both disks, limiting performance. Doubly distorted mirrors increase the number of physical writes per logical write from 2 to 3, but performs logical writes more efficiently. This reduces the cost of a random logical write to 1/3 of the cost of a read. Moreover, much of the write ...

6 File system aging—increasing the relevance of file system benchmarks

 Keith A. Smith, Margo I. Seltzer

June 1997 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1997 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '97**, Volume 25 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.81 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Benchmarks are important because they provide a means for users and researchers to characterize how their workloads will perform on different systems and different system architectures. The field of file system design is no different from other areas of research in this regard, and a variety of file system benchmarks are in use, representing a wide range of the different user workloads that may be run on a file system. A realistic benchmark,

however, is only one of the tools that is required in ...

7 Scale and performance in a distributed file system

 John H. Howard, Michael L. Kazar, Sherri G. Menees, David A. Nichols, M. Satyanarayanan,

Robert N. Sidebotham, Michael J. West

February 1988 **ACM Transactions on Computer Systems (TOCS)**, Volume 6 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(2.38 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The Andrew File System is a location-transparent distributed file system that will eventually span more than 5000 workstations at Carnegie Mellon University. Large scale affects performance and complicates system operation. In this paper we present observations of a prototype implementation, motivate changes in the areas of cache validation, server process structure, name translation, and low-level storage representation, and quantitatively demonstrate Andrews ability to scale gracefully. W ...

8 Ext3cow: a time-shifting file system for regulatory compliance

 Zachary Peterson, Randal Burns

May 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(443.01 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The ext3cow file system, built on the popular ext3 file system, provides an open-source file versioning and snapshot platform for compliance with the versioning and auditability requirements of recent electronic record retention legislation. Ext3cow provides a *time-shifting* interface that permits a real-time and continuous view of data in the past. Time-shifting does not pollute the file system namespace nor require snapshots to be mounted as a separate file system. Further, ext3cow is i ...

Keywords: Versioning file systems, copy-on-write

9 A high performance multi-structured file system design

 Keith Muller, Joseph Pasquale

September 1991 **ACM SIGOPS Operating Systems Review , Proceedings of the thirteenth ACM symposium on Operating systems principles SOSP '91**, Volume 25 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(1.40 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

File system I/O is increasingly becoming a performance bottleneck in large distributed computer systems. This is due to the increased file I/O demands of new applications, the inability of any single storage structure to respond to these demands, and the slow decline of, disk access times (latency and seek) relative to the rapid increase in CPU speeds, memory size, and network bandwidth. We present a *multi-structured file system* designed for high bandwidth I/O and fast response. Our design ...

10 IRON file systems

 Vijayan Prabhakaran, Lakshmi N. Bairavasundaram, Nitin Agrawal, Haryadi S. Gunawi,

Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(323.82 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Commodity file systems trust disks to either work or fail completely, yet modern disks exhibit more complex failure modes. We suggest a new *fail-partial failure model* for disks, which incorporates realistic localized faults such as latent sector errors and block corruption. We then develop and apply a novel *failure-policy fingerprinting* framework, to investigate how commodity file systems react to a range of more realistic disk failures. We classify their failure policies in a new ...

Keywords: IRON file systems, block corruption, disks, fail-partial failure model, fault tolerance, internal, latent sector errors, redundancy, reliability, storage

- 11 [Improving the performance of log-structured file systems with adaptive methods](#) 
-  Jeanna Neefe Matthews, Drew Roselli, Adam M. Costello, Randolph Y. Wang, Thomas E. Anderson
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5
- Publisher:** ACM Press
Full text available:  pdf(2.18 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

- 12 [Fast and secure distributed read-only file system](#) 
-  Kevin Fu, M. Frans Kaashoek, David Mazières
February 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 1
- Publisher:** ACM Press
Full text available:  pdf(317.54 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Internet users increasingly rely on publicly available data for everything from software installation to investment decisions. Unfortunately, the vast majority of public content on the Internet comes with no integrity or authenticity guarantees. This paper presents the self-certifying read-only file system, a content distribution system providing secure, scalable access to public, read-only data. The read-only file system makes the security of published content independent from that of the distri ...

Keywords: File systems, read-only, security

- 13 [Reimplementing the Cedar file system using logging and group commit](#) 
-  R. Hagmann
November 1987 **ACM SIGOPS Operating Systems Review , Proceedings of the eleventh ACM Symposium on Operating systems principles SOSP '87**, Volume 21 Issue 5
- Publisher:** ACM Press
Full text available:  pdf(775.94 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The workstation file system for the Cedar programming environment was modified to improve its robustness and performance. Previously, the file system used hardware-provided labels on disk blocks to increase robustness against hardware and software errors. The new system does not require hardware disk labels, yet is more robust than the old system. Recovery is rapid after a crash. The performance of operations on file system metadata, e.g., file creation or open, is greatly improved. < ...

- 14 [An end-to-end approach to globally scalable network storage](#) 
- Micah Beck, Terry Moore, James S. Plank
August 2002 **ACM SIGCOMM Computer Communication Review , Proceedings of the**



2002 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '02, Volume 32 Issue 4

Publisher: ACM Press

Full text available: [pdf\(286.82 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper discusses the application of end-to-end design principles, which are characteristic of the architecture of the Internet, to network storage. While putting storage into the network fabric may seem to contradict end-to-end arguments, we try to show not only that there is no contradiction, but also that adherence to such an approach is the key to achieving true scalability of shared network storage. After discussing end-to-end arguments with respect to several properties of network storage ...

Keywords: IBP, asynchronous communications, end-to-end design, exNode, internet backplane protocol, logistical networking, network storage, scalability, store and forward network, wide area storage

15 Fault tolerance under UNIX



Anita Borg, Wolfgang Blau, Wolfgang Graetsch, Ferdinand Herrmann, Wolfgang Oberle
January 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 1

Publisher: ACM Press

Full text available: [pdf\(1.97 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The initial design for a distributed, fault-tolerant version of UNIX based on three-way atomic message transmission was presented in an earlier paper [3]. The implementation effort then moved from Auragen Systems1 to Nixdorf Computer where it was completed. This paper describes the working system, now known as the TARGON/32. The original design left open questions in at least two areas: fault tolerance for server processes and recovery after a crash were brief ...

16 Replication in the harp file system



Barbara Liskov, Sanjay Ghemawat, Robert Gruber, Paul Johnson, Liuba Shrira
September 1991 **ACM SIGOPS Operating Systems Review , Proceedings of the thirteenth ACM symposium on Operating systems principles SOSP '91**, Volume 25 Issue 5

Publisher: ACM Press

Full text available: [pdf\(1.60 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes the design and implementation of the Harp file system. Harp is a replicated Unix file system accessible via the VFS interface. It provides highly available and reliable storage for files and guarantees that file operations are executed atomically in spite of concurrency and failures. It uses a novel variation of the primary copy replication technique that provides good performance because it allows us to trade disk accesses for network communication. Harp is intended to be used ...

17 X-RAY: A Non-Invasive Exclusive Caching Mechanism for RAIDs



Lakshmi N. Bairavasundaram, Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau
March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04**, Volume 32 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: [pdf\(250.59 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#)

RAID storage arrays often possess gigabytes of RAM for caching disk blocks. Currently,

most RAID systems use LRU or LRU-like policies to manage these caches. Since these arraycaches do not recognize the presence of file system buffer caches, they redundantly retain many of the same blocks as those cached by the file system, thereby wasting precious cache space. In this paper, we introduce X-RAY, an exclusive RAID array caching mechanism. X-RAY achieves a high degree of (but not perfect) exclusivity ...

18 Design and implementation of a configurable mixed-media file system

 Silvano Maffei

October 1994 **ACM SIGOPS Operating Systems Review**, Volume 28 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(333.42 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

In this paper we describe the design and implementation of a configurable mixed-media file system. The attribute *configurable* means that a file system serving a specific application area can be realized out of a library of reusable file system classes. The attribute *mixed-media* stands for the file system's ability to integrate different media types (RAM, harddisks, WORM optical disks, CDROMs, tape devices, RAIDs etc.) into a virtual storage, and making applications unaware of this ...

19 A cryptographic file system for UNIX

 Matt Blaze

December 1993 **Proceedings of the 1st ACM conference on Computer and communications security**

Publisher: ACM Press

Full text available:  [pdf\(955.62 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Although cryptographic techniques are playing an increasingly important role in modern computing system security, user-level tools for encrypting file data are cumbersome and suffer from a number of inherent vulnerabilities. The Cryptographic File System (CFS) pushes encryption services into the file system itself. CFS supports secure storage at the system level through a standard Unix file system interface to encrypted files. Users associate a cryptographic key with the directories ...

20 Recovery in the Calypso file system

 Murthy Devarakonda, Bill Kish, Ajay Mohindra

August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(318.88 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This article presents the design and implementation of the recovery scheme in Calypso. Calypso is a cluster-optimized, distributed file system for UNIX clusters. As in Sprite and AFS, Calypso servers are stateful and scale well to a large number of clients. The recovery scheme in Calypso is nondisruptive, meaning that open files remain open, client modified data are saved, and in-flight operations are properly handled across server recover. The scheme uses distributed state amount the client ...

Keywords: Calypso, cluster systems, distributed state, state reconstruction

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [!\[\]\(7377a3302f3d0fb3a834bf90f4594228_img.jpg\) Adobe Acrobat](#) [!\[\]\(280d20eb30edda888ea49b11139a3be7_img.jpg\) QuickTime](#) [!\[\]\(b5f97e7054ee78fe5f61d4f0345214fd_img.jpg\) Windows Media Player](#) [!\[\]\(e998c40a3119fe2e6fcad89de1fdcd94_img.jpg\) Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
[Search: The ACM Digital Library](#) [The Guide](#)

[THE ACM DIGITAL LIBRARY](#)
[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used
[inode file snapshot dataset source command metadata](#)

Found 14,301 of 169,166

Sort results by

[Save results to a Binder](#)
[Try an Advanced Search](#)

Display results

[Search Tips](#)
[Try this search in The ACM Guide](#)
 Open results in a new window

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 [Ext3cow: a time-shifting file system for regulatory compliance](#)

Zachary Peterson, Randal Burns

 May 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 2

Publisher: ACM Press

 Full text available: [pdf\(443.01 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The ext3cow file system, built on the popular ext3 file system, provides an open-source file versioning and snapshot platform for compliance with the versioning and auditability requirements of recent electronic record retention legislation. Ext3cow provides a *time-shifting* interface that permits a real-time and continuous view of data in the past. Time-shifting does not pollute the file system namespace nor require snapshots to be mounted as a separate file system. Further, ext3cow is i ...

Keywords: Versioning file systems, copy-on-write

2 [Fast detection of communication patterns in distributed executions](#)

Thomas Kunz, Michiel F. H. Seuren

 November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**
Publisher: IBM Press

 Full text available: [pdf\(4.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

3 [Services: ELF: an efficient log-structured flash file system for micro sensor nodes](#)

Hui Dai, Michael Neufeld, Richard Han

 November 2004 **Proceedings of the 2nd international conference on Embedded networked sensor systems**
Publisher: ACM Press

 Full text available: [pdf\(291.68 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An efficient and reliable file storage system is important to micro sensor nodes so that

data can be logged for later asynchronous delivery across a multi-hop wireless sensor network. Designing and implementing such a file system for a sensor node faces various challenges. Sensor nodes are highly resource constrained in terms of limited runtime memory, limited persistent storage, and finite energy. Also, the flash storage medium on sensor nodes differs in a variety of ways from the traditional ...

Keywords: eeprom, file system, flash, log structured, reliability, sensor

4 An algebraic approach to file synchronization

 Norman Ramsey, El'ad Csirmaz

September 2001 **ACM SIGSOFT Software Engineering Notes , Proceedings of the 8th European software engineering conference held jointly with 9th ACM SIGSOFT international symposium on Foundations of software engineering ESEC/FSE-9**, Volume 26 Issue 5

Publisher: ACM Press

Full text available:  pdf(301.78 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A *file synchronizer* restores consistency after multiple replicas of a filesystem have been changed independently. We present an algebra for reasoning about operations on filesystems and show that it is sound and complete with respect to a simple model. The algebra enables us to specify a file-synchronization algorithm that can be combined with several different conflict-resolution policies. By contrast, previous work builds the conflict-resolution policy into the specification, or worse, ...

5 Lineage retrieval for scientific data processing: a survey

 Rajendra Bose, James Frew

March 2005 **ACM Computing Surveys (CSUR)**, Volume 37 Issue 1

Publisher: ACM Press

Full text available:  pdf(728.75 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Scientific research relies as much on the dissemination and exchange of data sets as on the publication of conclusions. Accurately tracking the lineage (origin and subsequent processing history) of scientific data sets is thus imperative for the complete documentation of scientific work. Researchers are effectively prevented from determining, preserving, or providing the lineage of the computational data products they use and create, however, because of the lack of a definitive model for lineage ...

Keywords: Data lineage, audit, data provenance, scientific data, scientific workflow

6 Real-time shading

 Marc Olano, Kurt Akeley, John C. Hart, Wolfgang Heidrich, Michael McCool, Jason L. Mitchell, Randi Rost

August 2004 **Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04**

Publisher: ACM Press

Full text available:  pdf(7.39 MB) Additional Information: [full citation](#), [abstract](#)

Real-time procedural shading was once seen as a distant dream. When the first version of this course was offered four years ago, real-time shading was possible, but only with one-of-a-kind hardware or by combining the effects of tens to hundreds of rendering passes. Today, almost every new computer comes with graphics hardware capable of interactively executing shaders of thousands to tens of thousands of instructions. This course has been redesigned to address today's real-time shading capabili ...

7 Visualizing geospatial data

 Theresa Marie Rhyne, Alan MacEachern, Theresa-Marie Rhyne
August 2004 **Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04**

Publisher: ACM Press

Full text available:  pdf(13.99 MB) Additional Information: [full citation](#), [abstract](#)

This course reviews concepts and highlights new directions in GeoVisualization. We review four levels of integrating geospatial data and geographic information systems (GIS) with scientific and information visualization (VIS) methods. These include:

- Rudimentary: minimal data sharing between the GIS and Vis systems
- Operational: consistency of geospatial data
- Functional: transparent communication between the GIS and Vis systems
- Merged: one comprehensive toolkit environment

W ...

8 Speculative execution in a distributed file system

 Edmund B. Nightingale, Peter M. Chen, Jason Flinn
October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue 5

Publisher: ACM Press

Full text available:  pdf(305.54 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Speculator provides Linux kernel support for speculative execution. It allows multiple processes to share speculative state by tracking causal dependencies propagated through inter-process communication. It guarantees correct execution by preventing speculative processes from externalizing output, e.g., sending a network message or writing to the screen, until the speculations on which that output depends have proven to be correct. Speculator improves the performance of distributed file systems ...

Keywords: causality, distributed file systems, speculative execution

9 Routing design in operational networks: a look from the inside

 Geoffrey Xie, Jibin Zhan, David A. Maltz, Hui Zhang, Albert Greenberg, Gísli Hjálmtýsson
August 2004 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '04**, Volume 34 Issue 4

Publisher: ACM Press

Full text available:  pdf(372.95 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In any IP network, routing protocols provide the intelligence that takes a collection of physical links and transforms them into a network that enables packets to travel from one host to another. Though routing design is arguably the single most important design task for large IP networks, there has been very little systematic investigation into how routing protocols are actually used in production networks to implement the goals of network architects. We have developed a methodology for reverse ...

Keywords: network modeling, reverse engineering, routing design, static configuration analysis

10 Kernel korner: IBM's journaled filesystem

Steve Best, David Gordon, Ibrahim Haddad
January 2003 **Linux Journal**, Volume 2003 Issue 105

Publisher: Specialized Systems Consultants, Inc.

Full text available:  html(20.45 KB) Additional Information: [full citation](#), [index terms](#)

11 Soft updates: a solution to the metadata update problem in file systems

 Gregory R. Ganger, Marshall Kirk McKusick, Craig A. N. Soules, Yale N. Patt
May 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 2

Publisher: ACM Press

Full text available:  pdf(147.90 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Metadata updates, such as file creation and block allocation, have consistently been identified as a source of performance, integrity, security, and availability problems for file systems. Soft updates is an implementation technique for low-cost sequencing of fine-grained updates to write-back cache blocks. Using soft updates to track and enforce metadata update dependencies, a file system can safely use delayed writes for almost all file operations. This article describes soft ...

12 A Self-Organizing Storage Cluster for Parallel Data-Intensive Applications

Hong Tang, Aziz Gulbeden, Jingyu Zhou, William Strathern, Tao Yang, Lingkun Chu
November 2004 **Proceedings of the 2004 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society

Full text available:  pdf(330.26 KB) Additional Information: [full citation](#), [abstract](#)

Cluster-based storage systems are popular for data-intensive applications and it is desirable yet challenging to provide incremental expansion and high availability while achieving scalability and strong consistency. This paper presents the design and implementation of a self-organizing storage cluster called Sorrento, which targets data-intensive workload with highly parallel requests and low write-sharing patterns. Sorrento automatically adapts to storage node joins and departures, and the sys ...

13 Web3D in ocean science learning environments: virtual big beef creek

 Bruce Campbell, Paul Collins, Hunter Hadaway, Nick Hedley, Mark Stoermer
February 2002 **Proceeding of the seventh international conference on 3D Web technology**

Publisher: ACM Press

Full text available:  pdf(387.03 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The Virtual Reality Modeling Language (VRML), Java 3D software development packages, and World Wide Web (the Web) offer great potential for delivering three-dimensional, collaborative virtual environments to broad, on-line audiences. These capabilities have significant potential in ocean sciences, so a visualization environment was developed to explore these possibilities. The University of Washington's Virtual Big Beef Creek (VBBC) project has been continuously refined since its initial impleme ...

Keywords: VRML, interface paradigms, virtual environments, virtual geography, virtual worlds

14 Dynamic Metadata Management for Petabyte-Scale File Systems

Sage A. Weil, Kristal T. Pollack, Scott A. Brandt, Ethan L. Miller
November 2004 **Proceedings of the 2004 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society

Full text available:  pdf(175.04 KB) Additional Information: [full citation](#), [abstract](#)

In petabyte-scale distributed file systems that decouple read and write from metadata operations, behavior of the metadata server cluster will be critical to overall system performance and scalability. We present a dynamic subtree partitioning and adaptive metadata management system designed to efficiently manage hierarchical metadata

workloads that evolve over time. We examine the relative merits of our approach in the context of traditional workload partitioning strategies, and demonstrate the ...

- 15 [Papers: Information visualization: PhotoMesa: a zoomable image browser using quantum treemaps and bubblemaps](#)

 Benjamin B. Bederson

November 2001 **Proceedings of the 14th annual ACM symposium on User interface software and technology**

Publisher: ACM Press

Full text available:  pdf(1.34 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

PhotoMesa is a zoomable image browser that uses a novel treemap algorithm to present large numbers of images grouped by directory, or other available metadata. It uses a new interaction technique for zoomable user interfaces designed for novices and family use that makes it straightforward to navigate through the space of images, and impossible to get lost. PhotoMesa groups images using one of two new algorithms that lay out groups of objects in a 2D space-filling manner. *Quantum treemaps* ...

Keywords: Animation, Graphics, Image Browsers, Jazz, Treemaps, Zoomable User Interfaces (ZUIs)

- 16 [Frangipani: a scalable distributed file system](#)

 Chandramohan A. Thekkath, Timothy Mann, Edward K. Lee

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available:  pdf(2.20 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

- 17 [The Integrated Dictionary/Directory System](#)

 Frank W. Allen, Mary E. S. Loomis, Michael V. Mannino

June 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 2

Publisher: ACM Press

Full text available:  pdf(2.71 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

- 18 [IRON file systems](#)

 Vijayan Prabhakaran, Lakshmi N. Bairavasundaram, Nitin Agrawal, Haryadi S. Gunawi,

Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue 5

Publisher: ACM Press

Full text available:  pdf(323.82 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Commodity file systems trust disks to either work or fail completely, yet modern disks exhibit more complex failure modes. We suggest a new *fail-partial failure model* for disks, which incorporates realistic localized faults such as latent sector errors and block corruption. We then develop and apply a novel *failure-policy fingerprinting* framework, to investigate how commodity file systems react to a range of more realistic disk failures. We classify their failure policies in a new ...

Keywords: IRON file systems, block corruption, disks, fail-partial failure model, fault tolerance, internal, latent sector errors, redundancy, reliability, storage

- 19 [Embedded systems: applications, solutions and techniques \(EMBS\): A fast start-up technique for flash memory based computing systems](#) 

Keun Soo Yim, Jihong Kim, Kern Koh

March 2005 **Proceedings of the 2005 ACM symposium on Applied computing**

Publisher: ACM Press

Full text available:  pdf(324.29 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Flash memory based embedded computing systems are becoming increasingly prevalent. These systems typically have to provide an instant start-up time. However, we observe that mounting a file system for flash memory takes 1 to 25 seconds mainly depending on the flash capacity. Since the flash chip capacity is doubled in every year, this mounting time will soon become the most dominant reason of the delay of system start-up time. Therefore, in this paper, we present instant mounting techniques for ...

Keywords: fast booting, fast mounting, flash memory, metadata snapshot

- 20 [A Computational Database System for Generating Unstructured Hexahedral Meshes with Billions of Elements](#) 

Tiankai Tu, David R. O'Hallaron

November 2004 **Proceedings of the 2004 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society

Full text available:  pdf(222.13 KB) Additional Information: [full citation](#), [abstract](#)

For a large class of physical simulations with relatively simple geometries, unstructured octree-based hexahedral meshes provide a good compromise between adaptivity and simplicity. However, generating unstructured hexahedral meshes with over 1 billion elements remains a challenging task. We propose a database approach to solve this problem. Instead of merely storing generated meshes into conventional databases, we have developed a new kind of software system called Computational Database System ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

 **PORTAL**
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login
 Search: The ACM Digital Library The Guide



 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [inode disk address](#)

Found 11,284 of 169,166

Sort results by

 [Save results to a Binder](#)

[Try an Advanced Search](#)

Display results

 [Search Tips](#)
 [Open results in a new window](#)

[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale 

1 [4.2BSD and 4.3BSD as examples of the UNIX system](#)

 John S. Quarterman, Abraham Silberschatz, James L. Peterson
 December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(4.07 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This paper presents an in-depth examination of the 4.2 Berkeley Software Distribution, Virtual VAX-11 Version (4.2BSD), which is a version of the UNIX Time-Sharing System. There are notes throughout on 4.3BSD, the forthcoming system from the University of California at Berkeley. We trace the historical development of the UNIX system from its conception in 1969 until today, and describe the design principles that have guided this development. We then present the internal data structures and ...

2 [On-line data compression in a log-structured file system](#)

 Michael Burrows, Charles Jerian, Butler Lampson, Timothy Mann
 September 1992 **ACM SIGPLAN Notices, Proceedings of the fifth international conference on Architectural support for programming languages and operating systems ASPLOS-V**, Volume 27 Issue 9

Publisher: ACM Press

Full text available:  [pdf\(1.02 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

3 [Making Inodes Behave](#)

Clay Claiborne
 February 2001 **Linux Journal**

Publisher: Specialized Systems Consultants, Inc.

Full text available:  [html\(19.47 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Claiborne describes the difficulties he encountered while building Linux systems for General Dynamics.

4 [Deciding when to forget in the Elephant file system](#)

 Douglas S. Santry, Michael J. Feeley, Norman C. Hutchinson, Alistair C. Veitch, Ross W. Carton, Jacob Ofir
 December 1999 **ACM SIGOPS Operating Systems Review, Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP**

'99, Volume 33 Issue 5

Publisher: ACM PressFull text available:  pdf(1.61 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Modern file systems associate the deletion of a file with the immediate release of storage, and file writes with the irrevocable change of file contents. We argue that this behavior is a relic of the past, when disk storage was a scarce resource. Today, large cheap disks make it possible for the file system to protect valuable data from accidental delete or overwrite. This paper describes the design, implementation, and performance of the Elephant file system, which automatically retains all impo ...

5 [Research session: XML query processing #1: Efficient processing of XML path queries using the disk-based F&B Index](#)

Wei Wang, Haifeng Jiang, Hongzhi Wang, Xuemin Lin, Hongjun Lu, Jianzhong Li

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05****Publisher:** VLDB EndowmentFull text available:  pdf(371.23 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

With the proliferation of XML data and applications on the Internet, efficient XML query processing techniques are in great demand. Answering queries using XML indexes is a natural approach. A number of XML indexes have been proposed in the literature: among them, F&B Index is one powerful index as it is the smallest index that answers all twig queries. However, an F&B Index suffers from the following two problems: (1) it was originally proposed as a memory-based index while its size is usually ...

6 [The design and implementation of a log-structured file system](#)

 Mendel Rosenblum, John K. OusterhoutFebruary 1992 **ACM Transactions on Computer Systems (TOCS)**, Volume 10 Issue 1**Publisher:** ACM PressFull text available:  pdf(1.97 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This paper presents a new technique for disk storage management called a log-structured file system. A log-structured file system writes all modifications to disk sequentially in a log-like structure, thereby speeding up both file writing and crash recovery. The log is the only structure on disk; it contains indexing information so that files can be read back from the log efficiently. In order to maintain large free areas on disk for fast writing, we divide the log into

Keywords: Unix, disk storage management, fast crash recovery, file system organization, file system performance, high write performance, log-structured, logging

7 [The design and implementation of a log-structured file system](#)

 Mendel Rosenblum, John K. OusterhoutSeptember 1991 **ACM SIGOPS Operating Systems Review , Proceedings of the thirteenth ACM symposium on Operating systems principles SOSP '91**, Volume 25 Issue 5**Publisher:** ACM PressFull text available:  pdf(1.81 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents a new technique for disk storage management called a *log-structured file system*. A log-structured file system writes all modifications to disk sequentially in a log-like structure, thereby speeding up both file writing and crash recovery. The log is the only structure on disk; it contains indexing information so that files can be read back from the log efficiently. In order to maintain large free areas on disk for fast writing, we divide

the log into *segments* an ...

8 Improving the efficiency of UNIX buffer caches

 A. Braunstein, M. Riley, J. Wilkes
November 1989 **ACM SIGOPS Operating Systems Review , Proceedings of the twelfth ACM symposium on Operating systems principles SOSP '89**, Volume 23 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.46 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper reports on the effects of using hardware virtual memory assists in managing file buffer caches in UNIX. A controlled experimental environment was constructed from two systems whose only difference was that one of them (XMF) used the virtual memory hardware to assist file buffer cache search and retrieval. An extensive series of performance characterizations was used to study the effects of varying the buffer cache size (from 3 Megabytes to 70 MB); I/O transfer sizes (from ...

9 Improving storage system availability with D-GRAID

 Muthian Sivathanu, Vijayan Prabhakaran, Andrea C. Arpaci-Dusseau, Remzi H. Arpacidusseau

May 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available:  pdf(700.30 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present the design, implementation, and evaluation of D-GRAID, a gracefully degrading and quickly recovering RAID storage array. D-GRAID ensures that most files within the file system remain available even when an unexpectedly high number of faults occur. D-GRAID achieves high availability through aggressive replication of semantically critical data, and fault-isolated placement of logically related data. D-GRAID also recovers from failures quickly, restoring only live file system data to a h ...

Keywords: Block-based storage, Disk array, RAID, fault isolation, file systems, smart disks

10 Fault tolerance under UNIX

 Anita Borg, Wolfgang Blau, Wolfgang Graetsch, Ferdinand Herrmann, Wolfgang Oberle
January 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.97 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The initial design for a distributed, fault-tolerant version of UNIX based on three-way atomic message transmission was presented in an earlier paper [3]. The implementation effort then moved from Auragen Systems1 to Nixdorf Computer where it was completed. This paper describes the working system, now known as the TARGON/32. The original design left open questions in at least two areas: fault tolerance for server processes and recovery after a crash were brie ...

11 Frangipani: a scalable distributed file system

 Chandramohan A. Thekkath, Timothy Mann, Edward K. Lee
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available:  pdf(2.20 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 Techniques to increase disk access locality in the Minorca multimedia file system Chuanbao Wang, Vera Goebel, Thomas Plagemann October 1999 **Proceedings of the seventh ACM international conference on Multimedia (Part 2)****Publisher:** ACM PressFull text available:  pdf(483.67 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**Keywords:** access locality, continuous media, disk layout, disk seek, read ahead**13 The file system of an integrated local network** Paul J. Leach, Paul H. Levine, James A. Hamilton, Bernard L. Stumpf March 1985 **Proceedings of the 1985 ACM thirteenth annual conference on Computer Science****Publisher:** ACM PressFull text available:  pdf(1.78 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The distributed file system component of the DOMAIN system is described. The DOMAIN system is an architecture for networks of personal workstations and servers which creates an integrated distributed computing environment. The distinctive features of the file system include: objects addressed by unique identifiers (UIDs); transparent access to objects, regardless of their location in the network; the abstraction of a single level store for accessing all objects; and the layering of a networ ...

14 A high resolution disk I/O trace system Tao Huang, Teng Xu, Xianliang Lu October 2001 **ACM SIGOPS Operating Systems Review**, Volume 35 Issue 4**Publisher:** ACM PressFull text available:  pdf(377.26 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Disk access patterns are more important to understand as the gap between processor and disk performance is increasing. Obtaining disk I/O traces from real system is the first step to analyze disk access patterns, and tracing disk I/O is the first step to obtain disk I/O traces from real system. This paper implements a high resolution disk I/O trace system built into Linux, without adding noticeable processor load to the system. Each trace record contains the following content about a single phys ...

Keywords: Disk access pattern, Hard disk, Linux, Trace**15 Implementing global memory management in a workstation cluster** M. J. Feeley, W. E. Morgan, E. P. Pighin, A. R. Karlin, H. M. Levy, C. A. Thekkath December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95**, Volume 29 Issue 5**Publisher:** ACM PressFull text available:  pdf(1.52 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**16 ECOSystem: managing energy as a first class operating system resource**

Heng Zeng, Carla S. Ellis, Alvin R. Lebeck, Amin Vahdat

 October 2002 **ACM SIGPLAN Notices , ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , Proceedings of the 10th international conference on Architectural support for programming languages and operating systems ASPLOS-X**, Volume 37 , 30 , 36 Issue 10 , 5 , 5

Publisher: ACM Press

Full text available:  pdf(1.17 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Energy consumption has recently been widely recognized as a major challenge of computer systems design. This paper explores how to support energy as a first-class operating system resource. Energy, because of its global system nature, presents challenges beyond those of conventional resource management. To meet these challenges we propose the Currency Model that unifies energy accounting over diverse hardware components and enables fair allocation of available energy among applications. Our par ...

17 A generalized text editor

 Christopher W. Fraser

March 1980 **Communications of the ACM**, Volume 23 Issue 3

Publisher: ACM Press

Full text available:  pdf(512.06 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Text is not the only data that needs editing; for example, file deletion utilities edit directories. If all "editors" used the same command language, they would be easier to learn, remember, and code. This paper describes a generalized editor that edits text, directories, binary core images, and certain operating system data with a single user interface.

Keywords: CRT, command language, editor, text

18 Probing the black box: Transforming policies into mechanisms with infokernel

 Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, Nathan C. Burnett, Timothy E. Denehy, Thomas J. Engle, Haryadi S. Gunawi, James A. Nugent, Florentina I. Popovici
October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**

Publisher: ACM Press

Full text available:  pdf(365.12 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We describe an evolutionary path that allows operating systems to be used in a more flexible and appropriate manner by higher-level services. An infokernel exposes key pieces of information about its algorithms and internal state; thus, its default policies become mechanisms, which can be controlled from user-level. We have implemented two prototype infokernels based on the linuxtwofour and netbsdver kernels, called infolinix and infobsd, respectively. The infokernels export key abstractions as ...

Keywords: information, mechanism, policy

19 Implementing cooperative prefetching and caching in a globally-managed memory system

 Geoffrey M. Voelker, Eric J. Anderson, Tracy Kimbrel, Michael J. Feeley, Jeffrey S. Chase, Anna R. Karlin, Henry M. Levy
June 1998 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1998 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems SIGMETRICS '98/PERFORMANCE '98**, Volume 26 Issue 1

Publisher: ACM Press

Full text available: [pdf\(1.66 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents *cooperative prefetching and caching* --- the use of network-wide global resources (memories, CPUs, and disks) to support prefetching and caching in the presence of hints of future demands. Cooperative prefetching and caching effectively unites disk-latency reduction techniques from three lines of research: prefetching algorithms, cluster-wide memory management, and parallel I/O. When used together, these techniques greatly increase the power of prefetching relative to a ...

20 [Measuring file access patterns in UNIX](#)



Irene Hu

August 1986 **ACM SIGMETRICS Performance Evaluation Review**, Volume 14 Issue 2

Publisher: ACM Press

Full text available: [pdf\(473.91 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

UNIX is a disk-based operating system, where only the system kernel is always memory-resident. A combination of small block size, limited read-ahead and numerous seeks can severely limit the file system throughput. This paper presents a tool to study the file access patterns. Information derived from the data collected can be used to determine the optimal disk block size and also to improve the block placement strategy. The tool is a software monitor, installed at the device driver level, and tr ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#)^{New!} [more »](#)

no metadata store file initial snapshot

[Advanced Search](#)[Preferences](#)**Web**Results 1 - 10 of about 244,000 for **no metadata store file initial snapshot**. (0.28 seconds)**Enhancing Merge Replication Performance (Replication (SQL Server))**Creating a ROWGUIDCOL column prior to generating the **initial snapshot** ...on the Distributor because **no meta data** has to be **stored** about the Subscriber. ...msdn.microsoft.com/library/en-us/replsql/replperf_5kx1.asp - 27k - [Cached](#) - [Similar pages](#)**How Merge Replication Works (Replication (SQL Server))**The **Snapshot Agent** prepares **snapshot files** containing schema and data of published tables, **stores** the files in the **snapshot** folder, ...msdn.microsoft.com/library/en-us/replsql/repltypes_30z7.asp - 34k - [Cached](#) - [Similar pages](#)[[More results from msdn.microsoft.com](#)]**[PDF] THE LAST WORD IN FILE SYSTEMS**File Format: PDF/Adobe Acrobat - [View as HTML](#)100% dynamic **metadata**. • **No** limits on **files**, directory entries, etc. ...All **storage** is shared – **no** wasted space, **no** wasted bandwidth ...www.opensolaris.org/os/community/zfs/docs/zfs_last.pdf - [Similar pages](#)**[PDF] Focusing on Snapshots**

File Format: PDF/Adobe Acrobat

backup **file** by **file**. Snapshots are the equivalent of raw backups, the. **snapshot**hardware has **no** facility for handling **meta-data**. It is the backup ...www.naspa.com/PDF/2004/1204/N0412006.pdf - [Similar pages](#)**[PDF] A Storage Networking Appliance**File Format: PDF/Adobe Acrobat - [View as HTML](#)consumes **no** disk space until **files** referenced by a **Snapshot** are deleted or modified.... Like Episode, WAFL uses **files** to **store** **meta-data**. ...www.netapp.com/tech_library/3001.html - [Similar pages](#)**[PDF] STREAMLINING ORACLE E-BUSINESS SUITE 11 i OPERATIONS AND MIGRATION ...**File Format: PDF/Adobe Acrobat - [View as HTML](#)A **Snapshot file** can be. created in just a few seconds and uses **no** ... Once a **Snapshot file** is created, the **storage** appliance can back up the database ...www.netapp.com/tech_library/3310.html - [Similar pages](#)**[PDF] Meta-data Snapshotting: A Simple Mechanism for File System Consistency**File Format: PDF/Adobe Acrobat - [View as HTML](#)altered **meta-data** blocks, representing such **files**, in a **snapshot**. The **snapshot**will be scheduled to be written to stable **storage** ...www.research.ibm.com/K42/papers/snapi03.pdf - [Similar pages](#)**snapfs**... is that if all the filesystem **metadata** is flushed before making the **snapshot**,... tiny **files**. ext2fs **stores** short symlinks within the inode itself. ...sourcefrog.net/projects/snapfs/ - 18k - [Cached](#) - [Similar pages](#)**Network File System Version 4 (nfsv4) Charter**

2.7.9 Network File System Version 4 (nfsv4). NOTE: This charter is a **snapshot** of the 62nd IETF Meeting in Minneapolis, MN USA. It may now be out-of-date. ...
www3.ietf.org/proceedings/05mar/nfsv4.html - 22k - [Cached](#) - [Similar pages](#)

[PDF] **Richer File System Metadata Using Links and Attributes**

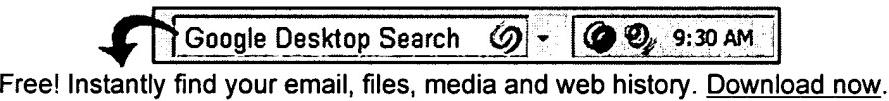
File Format: PDF/Adobe Acrobat - [View as HTML](#)

employ far richer **metadata** structures. File systems design-**ers** will no longer be ... the underlying Linux **file system** for data **storage**. Initial ...

www.cs.ucsc.edu/~nikhil/Papers/msst05richer.pdf - [Similar pages](#)

Try your search again on [Google Book Search](#)

Gooooooooogle ►
Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)



[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2005 Google

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#)^{New!} [more »](#)

inode snapshot file system

[Search](#)
[Advanced Search](#)
[Preferences](#)
WebResults 1 - 10 of about 129,000 for **inode snapshot file system**. (0.29 seconds)**[PDF] File System Design for an NFS File Server Appliance**File Format: PDF/Adobe Acrobat - [View as HTML](#)

as the root **inode** represents the active **file system**. When the **Snapshot inode** is created, it points to exactly the same disk blocks as the root **inode**, ...

www.netapp.com/tech_library/3002.html - [Similar pages](#)

Sponsored Links

File snapshot

Create data snapshots & allow the backup of open & in-use files.

www.stbernard.com
Snapshot File Systems

When an **inode** update or a write changes the data in block n of the snapped **file system**, the old data is first read and copied to the **snapshot** before the ...

docs.hp.com/en/B3929-90011/ch04s02.html - 28k - [Cached](#) - [Similar pages](#)
Using a Snapshot File System for Backup

The **file system** may have some extended **inode** operations that must be completed, though there should be no other changes. Since the **snapshot file system** is ...

docs.hp.com/en/B3929-90011/ch04s03.html - 31k - [Cached](#) - [Similar pages](#)
[\[More results from docs.hp.com \]](#)
LFS: A Log Structured File System for Linux that Supports Snapshots

Traditional **file systems** like ext2 usually write **inode** blocks at a fixed place ... No current **Linux file system** supports **snapshots** and implementing a **file** ...

logfs.sourceforge.net/ - 10k - [Cached](#) - [Similar pages](#)
[PDF] A Log Structured File System with SnapshotsFile Format: PDF/Adobe Acrobat - [View as HTML](#)

Figure 4: **Snapshots**, shown in (4). Note that the **inode** map itself is a **file** in our **file system** and modification of **inode** map results ...

www.eecs.umich.edu/~ppadala/soc/lfs.pdf - [Similar pages](#)
The snapshot archive available at this web site contains snapshot ...

For example, 15Oct94/sd2g contains the **snapshot** of the glan5 **file system**. ... It describes the **file** corresponding to the **inode**. number = **inode** number gen ...

www.eecs.harvard.edu/~keith/snapshots/description - 4k - [Cached](#) - [Similar pages](#)
GPFS V2.3 Administration and Programming Reference ...

gpfs_open_inodescan() Subroutine. Name. gpfs_open_inodescan() - Opens an **inode** scan of a **file system** or **snapshot**. Library. GPFS Library (libgpfs.a for AIX, ...
publib.boulder.ibm.com/infocenter/clresctr/

topic/com.ibm.cluster.gpfs.doc/gpfs23/bl1adm10/bl1adm10161.html - 11k -
[Cached](#) - [Similar pages](#)
GPFS V2.3 Administration and Programming Reference - Managing GPFS ...

Snapshots are exact copies of changed data in the active **files** and directories of a **file system** with the exception of the **inode** number. ...

publib.boulder.ibm.com/infocenter/clresctr/
topic/com.ibm.cluster.gpfs.doc/gpfs23/bl1adm10/bl1adm1030.html - 18k -
[Cached](#) - [Similar pages](#)

Network storage file systems for the enterprise

All **file systems** have **inodes**, or something similar. Until recently **file system** designers

studiously attempted to minimize **inode** size (typically 64 or 128 ...
www.traakan.com/products.html - 20k - [Cached](#) - [Similar pages](#)

[PDF] [**An Efficient Snapshot Technique for Ext3 File System in Linux 2.6**](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

In this paper, we develop a **file system-based snapshot** for the ext3 file ... and the struct **inode** defines all information needed. by the **file system** to ...

www.linuxdevices.com/files/rtlws-2005/SeungjunShim.pdf - [Similar pages](#)

Try your search again on [Google Book Search](#)

Gooooooooooooogle ►

Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)



Free! Instantly find your email, files, media and web history. [Download now.](#)

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2005 Google

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#)^{New!} [more »](#)

inode snapshot file system

[Search](#)[Advanced Search](#)
[Preferences](#)**Web**Results 11 - 20 of about 129,000 for **inode snapshot file system**. (0.09 seconds)**snapfs**

If the **system** wants to write to a data block, **inode**, or directory entry used by a previous **snapshot**, then the **filesystem** transparently allocates a block ...
sourcefrog.net/projects/snapfs/ - 18k - [Cached](#) - [Similar pages](#)

Sponsored Links**File snapshot**

15-day free trial. Back up open files to protect your business.
www.stbernard.com

Network storage file systems for the enterprise

... disk usage by integrating **Snapshot** features directly into their **file systems** ... Additionally 50% of all **files** are small enough to be included in **inode** ...
www.traakan.com/products_features.html - 20k - [Cached](#) - [Similar pages](#)

GNU tar

If you don't use --incremental (-G), the **file system** will probably fill up with **files** that shouldn't exist any more. --listed-incremental= **snapshot-file** ...
www.delorie.com/gnu/docs/tar/tar_78.html - 10k - [Cached](#) - [Similar pages](#)

inode.5

... list of **snapshot inode** numbers */ /* back to stuff that has been around a ... The root **inode** is the root of the **file system**. **Inode** 0 can't be used for ...
www.daemon-systems.org/man/inode.5.html - 17k - [Cached](#) - [Similar pages](#)

Glossary

inode. A unique identifier for each **file** within a **file system** which also ... A **file system** whose exact image has been used to create a **snapshot file system** ...
uw713doc.sco.com/ODM_FSSadmin/fssag-12.html - 14k - [Cached](#) - [Similar pages](#)

Introspective: What's a Snapshot Backup?

Performing a restore merely requires finding the **filesystem snapshot** for the ... Each unique **file** on the **filesystem** has a magic number, called an **inode**, ...
edseek.com/~jasonb/articles/dirvish_backup/snapshot.html - 10k - [Cached](#) - [Similar pages](#)

The ext3cow File System

... enhanced version of ext3 that allows one to take a **snapshot** of one's **file system**, ... For every version of a **file** that exists, an **inode** must exist to ...
www.ext3cow.com/ - 13k - [Cached](#) - [Similar pages](#)

[PDF] Journalized File System (JFS) for Linux UT, Texas April 25, 2003

File Format: PDF/Adobe Acrobat - [View as HTML](#)

File layout (**inode** containing the root of a B+tree which ... Setup the volume to use as the **snapshot**. Stop the **File System** operations (VFS operation) ...
www.cs.utexas.edu/users/ygz/378-03S/IBM-jfs.pdf - [Similar pages](#)

Linux Kernel 2.4 Internals: Virtual Filesystem (VFS)

This brings us from **inode** cache back to the **filesystem** code - remember that ... If this was an anonymous **inode** (NetApp .**snapshot**) then we delete it from the ...
www.faqs.org/docs/kernel_2_4/lki-3.html - 62k - [Cached](#) - [Similar pages](#)

Re: how to use LVM snapshot with ext3 - VFS lock patch applicability

... to copy the original >superblock, **inode** tables, etc to the **snapshot** once. ... I would like to make a **snapshot** of my **filesystem** while the application is ...

<https://listman.redhat.com/archives/ext3-users/2003-June/msg00025.html> - 8k -
[Cached](#) - [Similar pages](#)

◀ Gooooooooooooogle ▶
Result Page: [Previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [11](#) [Next](#)

[Search within results](#) | [Language Tools](#) | [Search Tips](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2005 Google

[Sign in](#)

Web Images Groups News Froogle Local^{New!} more »
 ditto address snapshot files modify metadata store commands Advanced Search Preferences

Web Results 1 - 10 of about 489 for ditto address snapshot files modify metadata store commands . (0.37

Generating data set of the first file system by determining a set ...

The data stored within files in a file system have associated metadata to describe the data and ... If all more recent snapshots contain a "ditto" address, ...

www.freepatentsonline.com/6959310.html - 179k - [Cached](#) - [Similar pages](#)

Generating data set of the first file system by determining a set ...

The data stored within files in a file system have associated metadata to ... If all more recent snapshots contain a "ditto" address, the lack of a more ...

www.patentstorm.us/patents/6959310.html - 176k - [Cached](#) - [Similar pages](#)

[PDF] Mac OS X Enterprise Application Management Best Practices

File Format: PDF/Adobe Acrobat - [View as HTML](#)

ditto is a command line tool which supports copying files and folders ... Mac OS X. Many applications under Mac OS X store metadata in a "resource fork", ...

macenterprise.org/dm/documents/20041020-937_bestpractices.pdf - [Similar pages](#)

IBM Globalization - Terminology

Virtual images are stored in the integrated file system and can be in either ...

Metadata about these resources resides in a directory on the file system; ...

www-306.ibm.com/software/globalization/terminology/vw.html - 145k - [Cached](#) - [Similar pages](#)

PostgreSQL

The dump (snapshot) file has just the ASCII SQL commands that specify the table contents. ... Dump the address book database to a dump file, including OIDs. ...

linuxmafia.com/faq/Apps/postgresql.html - 42k - [Cached](#) - [Similar pages](#)

Planet Sun (solaris)

To create a snapshot of our ZFS filesystem, we can use the command: ... You'll also need to modify your environment file to enable non-debug builds. ...

planetsun.org/filter/solaris/ - 101k - Dec 20, 2005 - [Cached](#) - [Similar pages](#)

[PDF] Mainframe Appliance for Storage

File Format: PDF/Adobe Acrobat - [View as HTML](#)

... updated/changed in order to support "snapshot" backups ... or Microsoft's Common Internet File System (CIFS ... the gigabit Ethernet interface with an IP address. ...

www.bustech.com/support/downloads/mas/40-02114-B0-001.pdf - Supplemental Result - [Similar pages](#)

ditto: Mac OS X

Anyway, I'm planning to make it use all system commands (mostly ditto) ...

Nothing in a Mac OS X 10.4 could be as cool as a metadata driven file system! ...

sean.typepad.com/ditto/mac_os_x/ - 348k - [Cached](#) - [Similar pages](#)

Subversion Dev: [PATCH book] Use of <quote> over "

-106,14 +106,15 @@ <firstterm>lock-modify-unlock</firstterm> model to address

... commit</command> of the file will fail with an "out-of-date" - error. ...

svn.haxx.se/dev/archive-2003-03/1274.shtml - 88k - [Cached](#) - [Similar pages](#)

[PDF] Integrity Checking For Process Hardening

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Other invalid **command**. Invalid shell **command**. File name binding. Directory confinement.

File name collision. Other **metadata**. Format string. Buffer overflow ...

https://www.cerias.purdue.edu/tools_and_resources/bibtex_archive/archive/Lhee.pdf - [Similar pages](#)

Try your search again on [Google Book Search](#)

Gooooooooooooogle ►

Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

Free! Get the Google Toolbar. [Download Now](#) - [About Toolbar](#)



[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2005 Google

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#)^{New!} [more »](#)

source file initial snapshot

[Search](#)
[Advanced Search](#)
[Preferences](#)
WebResults 21 - 30 of about 1,570,000 for **source file initial snapshot**. (0.09 seconds)**procmail history file (snapshot)**

It is derived from the HISTORY file that is included in source distributions.

The latest **snapshot** can be found at ...[www.procmail.org/procmail-history.html](#) - 33k - [Cached](#) - [Similar pages](#)**CHAPTER 5 RootCause Demo**Click on the pi_demo.cpp **source file** node, then right-click to see the Trace ...In this small demo, a **snapshot** is not really necessary since it doesn't ...[www.ocsystems.com/user_guide/ rootcause/win/html/rcc-8.htm](#) - 27k - [Cached](#) - [Similar pages](#)**C++ Development Environment - ugu**The PCH file contains a **snapshot** of all the code preceding the header stop point, typically the first token in the primary **source file** that does not belong ...[h30097.www3.hp.com/cplus/ugu_pch.html](#) - 35k - [Cached](#) - [Similar pages](#)**ScatterWeb.APP: Applications/PictureIt!/ESB/src/ScatterWeb.Camera ...**

00006 00007 Redistribution and use in source and binary forms, with or without

00008 modification, ... 00054 #define CAMERA_SYNCH (0x0D) // Initial Camera. ...

[www.inf.fu-berlin.de/inst/ag-tech/scatterweb_net/ ssi/docAPP/ESB_2src_2ScatterWeb_8Camera_8c-source.html](#) - 68k - [Cached](#) - [Similar pages](#)**New Repackage by Monitor and Snapshot**

It is a good idea to do an initial scan of a clean machine which can be used ...

It can be a good idea to copy all **source file** to a repository so that the ...[www.masaisolutions.com/help/installer/ RepackagingWizard/repackagebysnapshot.htm](#) - 8k - [Cached](#) - [Similar pages](#)**Comeau C++ 4.0 Pre-Release User Documentation: Precompiled Headers**The PCH file will contain a **snapshot** of all the code preceding the ... The initial sequence of preprocessing directives from the primary **source file**, ...[www.comeaucomputing.com/4.0/docs/userman/pch.html](#) - 17k - [Cached](#) - [Similar pages](#)**How to Use File Choosers**Try to use all of the controls on the file chooser. In the **source file** ...

You can specify the file chooser's initial directory using one of JFileChooser's ...

[java.sun.com/docs/books/tutorial/ uiswing/components/filechooser.html](#) - 49k - [Cached](#) - [Similar pages](#)**Structured Source Editing Evaluation Guide**The code in this initial proposed contribution is a **snapshot** "midstream" of our move to ... JSP files also participate in the moving of Java **source files**. ...[www.eclipse.org/webtools/initial-contribution/ IBM/evalGuides/SSEEval.html?p=1](#) - 35k - [Cached](#) - [Similar pages](#)**SnapshotCM Keyword Expansions**

\$Source\$, The full pathname of the file. \$Source: /path/file \$. \$Snapshot\$,

The full **snapshot** path for the **snapshot** used to do the check out. ...[www.truebluesoftware.com/docs/Keywords.html](#) - 17k - [Cached](#) - [Similar pages](#)**Re: gEDA: Need help resolving a scheduling contradiction**

... 3 3 Writing initial simulation **snapshot**: worklib.test:v Loading **snapshot** ...
Done ncsim> **source** /usr/ldv/tools/inca/files/ncsimrc ncsim> run FAILED ab=0 ...
archives.seul.org/geda/dev/Jun-2005/msg00038.html - 10k - Cached - Similar pages

◀ Gooooooooooooogle ▶

Result Page: **Previous** 1 2 3 4 5 6 7 8 9 10 11 12 **Next**

[Search within results](#) | [Language Tools](#) | [Search Tips](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2005 Google

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 [Search Session History](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Thu, 22 Dec 2005, 3:49:28 PM EST

Edit an existing query or
compose a new query in the
Search Query Display.

Search Query Display

Select a search number (#)
to:

- Add a query to the Search Query Display
 - Combine search queries using AND, OR, or NOT
 - Delete a search
 - Run a search
- Recent Search Queries**
- #1 (metadata<in>metadata)
#2 ((metadata<in>metadata) <and> (inode<in>metadata))
#3 ((metadata<in>metadata) <and> (inode<in>metadata))
#4 ((dataset<in>metadata) <and> (ditto<in>metadata))<and>
(address<in>metadata)
#5 ((copy<in>metadata) <and> (snapshot<in>metadata))
#6 ((inode<in>metadata) <and> (disk<in>metadata))<and>
(address<in>metadata)
#7 ((inode<in>metadata) <and> (unix<in>metadata))<and>
(file<in>metadata)
#8 ((file<in>metadata) <and> (system<in>metadata))<and>
(inode<in>metadata)
#9 ((command<in>metadata) <and> (modify<in>metadata))
<and> (inode<in>metadata)
#10 ((command<in>metadata) <and> (metadata<in>metadata))
<and> (file<in>metadata)
#11 ((creating<in>metadata) <and> (snapshot<in>metadata))
<and> (file<in>metadata)
#12 ((store<in>metadata) <and> (snapshot<in>metadata))<and>
(meta<in>metadata)
#13 ((store<in>metadata) <and> (snapshot<in>metadata))<and>
(meta<in>metadata)



[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |
Welcome United States Patent and Trademark Office

Search Results

BROWSE

SEARCH

IEEE XPLOR GUIDE

Results for "((copy<in>metadata) <and> (snapshot<in>metadata))"

[e-mail](#)

Your search matched 14 of 1286976 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance in Descending order**.

» **Search Options**

[View Session History](#)

[New Search](#)

Modify Search

((copy<in>metadata) <and> (snapshot<in>metadata)))

Check to search only within this results set

Display Format: **Citation** **Citation & Abstract**

» **Key**

IEEE JNL IEEE Journal or Magazine

Select Article Information

IEE JNL IEE Journal or Magazine

Select Article Information

IEEE CNF IEEE Conference Proceeding

- 1. **An intelligent page store for concurrent transaction and query processing**
Dias, D.M.; Goyal, A.; Parr, F.N.;
Research Issues on Data Engineering, 1992: Transaction and Query Processing International Workshop on
2-3 Feb. 1992 Page(s):12 - 19
Digital Object Identifier 10.1109/RIDE.1992.227429

[AbstractPlus](#) | [Full Text: PDF\(748 KB\)](#) **IEEE CNF**

IEEE STD IEEE Standard

- 2. **Parallel generation of base relation snapshots for materialized view main warehouse environment**
Saeki, S.; Bhalla, S.; Hasegawa, M.;
Parallel Processing Workshops, 2002. Proceedings. International Conference : 18-21 Aug. 2002 Page(s):383 - 390
Digital Object Identifier 10.1109/ICPPW.2002.1039755

[AbstractPlus](#) | [Full Text: PDF\(285 KB\)](#) **IEEE CNF**

- 3. **Unitary ESPRIT: how to obtain increased estimation accuracy with a reduced computational burden**
Haardt, M.; Nossek, J.A.;
Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]
Volume 43, Issue 5, May 1995 Page(s):1232 - 1242
Digital Object Identifier 10.1109/78.382406

[AbstractPlus](#) | [Full Text: PDF\(888 KB\)](#) **IEEE JNL**

- 4. **Error handling for the CDF online silicon vertex tracker**
Bari, M.; Belforte, S.; Cerri, A.; Dell'Orso, M.; Donati, S.; Galeotti, S.; Giannetti, P.; Punzi, G.; Ristori, L.; Spinella, F.; Zanetti, A.M.;
Nuclear Science, IEEE Transactions on
Volume 48, Issue 4, Part 1, Aug. 2001 Page(s):1286 - 1289
Digital Object Identifier 10.1109/23.958767

[AbstractPlus](#) | [References](#) | [Full Text: PDF\(73 KB\)](#) **IEEE JNL**

- 5. **Postgres-R(SI): Combining Replica Control with Concurrency Control Based Isolation**
Shuqing Wu; Kemme, B.;
Data Engineering, 2005. ICDE 2005. Proceedings. 21st International Conference : 05-08 April 2005 Page(s):422 - 433

Digital Object Identifier 10.1109/ICDE.2005.108

[AbstractPlus](#) | Full Text: [PDF\(328 KB\)](#) IEEE CNF

- 6. Percolation search in power law networks: making unstructured peer-to-peer scalable**
Sarshar, N.; Boykin, P.O.; Roychowdhury, V.P.;
Peer-to-Peer Computing, 2004. Proceedings. Proceedings. Fourth International Conference on
25-27 Aug. 2004 Page(s):2 - 9
Digital Object Identifier 10.1109/PTP.2004.1334925
[AbstractPlus](#) | Full Text: [PDF\(321 KB\)](#) IEEE CNF
- 7. Computation scrapbooks of Emacs Lisp runtime state**
Potter, R.;
Human-Centric Computing Languages and Environments, 2001. Proceedings on
5-7 Sept. 2001 Page(s):236 - 237
Digital Object Identifier 10.1109/HCC.2001.995267
[AbstractPlus](#) | Full Text: [PDF\(243 KB\)](#) IEEE CNF
- 8. Who links to whom: mining linkage between Web sites**
Bharat, K.; Bay-Wei Chang; Henzinger, M.; Ruhl, M.;
Data Mining, 2001. ICDM 2001, Proceedings IEEE International Conference on
29 Nov.-2 Dec. 2001 Page(s):51 - 58
Digital Object Identifier 10.1109/ICDM.2001.989500
[AbstractPlus](#) | Full Text: [PDF\(775 KB\)](#) IEEE CNF
- 9. Error handling for the CDF Silicon Vertex Tracker**
Belforte, S.; Cerri, A.; Dell'Orso, M.; Donati, S.; Galeotti, S.; Giannetti, P.; Morselli, P.; Ristori, L.; Spinella, F.; Zanetti, A.M.;
Nuclear Science Symposium Conference Record, 2000 IEEE
Volume 2, 15-20 Oct. 2000 Page(s):12/74 - 12/77 vol.2
Digital Object Identifier 10.1109/NSSMIC.2000.949945
[AbstractPlus](#) | Full Text: [PDF\(344 KB\)](#) IEEE CNF
- 10. Wait-free snapshots in real-time systems: algorithms and performance**
Ermedahl, A.; Hansson, H.; Papatriantafilou, M.; Tsigas, P.;
Real-Time Computing Systems and Applications, 1998. Proceedings. Fifth International Conference on
27-29 Oct. 1998 Page(s):257 - 266
Digital Object Identifier 10.1109/RTCSA.1998.726426
[AbstractPlus](#) | Full Text: [PDF\(152 KB\)](#) IEEE CNF
- 11. A locking protocol for multilevel secure databases using two committed phases**
Pal, S.;
Computer Assurance, 1995. COMPASS '95. 'Systems Integrity, Software Safe Security'. Proceedings of the Tenth Annual Conference on
25-29 June 1995 Page(s):197 - 210
Digital Object Identifier 10.1109/CMPASS.1995.521899
[AbstractPlus](#) | Full Text: [PDF\(1132 KB\)](#) IEEE CNF
- 12. Wait depth limited concurrency control**
Franceszek, P.A.; Robinson, J.T.; Thomasian, A.;
Data Engineering, 1991. Proceedings. Seventh International Conference on
8-12 April 1991 Page(s):92 - 101
Digital Object Identifier 10.1109/ICDE.1991.131456
[AbstractPlus](#) | Full Text: [PDF\(800 KB\)](#) IEEE CNF
- 13. A two snapshot algorithm for concurrency control in multi-level secure databases**

- Ammann, P.; Jaeckle, F.; Jajodia, S.;
Research in Security and Privacy, 1992. Proceedings., 1992 IEEE Computer Society Symposium on
4-6 May 1992 Page(s):204 - 215
Digital Object Identifier 10.1109/RISP.1992.213260
[AbstractPlus](#) | Full Text: [PDF\(912 KB\)](#) IEEE CNF
14. Scheduling the allocation of data fragments in a distributed database environment using a machine learning approach
Chaturvedi, A.R.; Choubey, A.K.; Jinsheng Roan;
Engineering Management, IEEE Transactions on
Volume 41, Issue 2, May 1994 Page(s):194 - 207
Digital Object Identifier 10.1109/17.293386
[AbstractPlus](#) | Full Text: [PDF\(1148 KB\)](#) IEEE JNL

[View All Selected Results](#)

[Help](#) [Contact Us](#) [Privacy & Terms](#)

© Copyright 2005 IEEE - All Rights Reserved





Home | Login | Logout | Access Information | Alerts |
Welcome United States Patent and Trademark Office

Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "((command<in>metadata) <and> (metadata<in>metadata))<and> (file<in...>"

[e-mail](#)

Your search matched 3 of 1286976 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.

» **Search Options**

[View Session History](#)

[New Search](#)

Modify Search

((command<in>metadata) <and> (metadata<in>metadata))<and> (file<in...>) [»](#)

Check to search only within this results set

Display Format: Citation Citation & Abstract

» **Key**

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

Select Article Information

- 1. **Grid Datafarm architecture for petascale data intensive computing**
Tatebe, O.; Morita, Y.; Matsuoka, S.; Soda, N.; Sekiguchi, S.;
Cluster Computing and the Grid 2nd IEEE/ACM International Symposium CCG
21-24 May 2002 Page(s):92 - 100
Digital Object Identifier 10.1109/CCGRID.2002.1017117
[AbstractPlus](#) | Full Text: [PDF\(276 KB\)](#) IEEE CNF
- 2. **I want my FTP: bits on demand**
Khare, R.;
Internet Computing, IEEE
Volume 2, Issue 4, July-Aug. 1998 Page(s):88 - 91
Digital Object Identifier 10.1109/4236.707811
[AbstractPlus](#) | Full Text: [PDF\(88 KB\)](#) IEEE JNL
- 3. **Grid Datafarm Architecture for Petascale Data Intensive Computing**
Tatebe, O.; Morita, Y.; Matsuoka, S.; Soda, N.; Sekiguchi, S.;
Cluster Computing and the Grid, 2002. 2nd IEEE/ACM International Symposium
21-24 May 2002 Page(s):102 - 102
Digital Object Identifier 10.1109/CCGRID.2002.1017117
[AbstractPlus](#) | Full Text: [PDF\(152 KB\)](#) IEEE CNF

[Advanced Search](#) | [Help](#) | [Logout](#) | [Feedback](#)

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2005 IEEE –

Indexed by
Inspec


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 Search Results**BROWSE****SEARCH****IEEE Xplore GUIDE**

Results for "(store<in>metadata) <and> (snapshot<in>metadata))<and> (meta<in>metadata)"

 [e-mail](#)

Your search matched 24 of 1286976 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance in Descending order**.**» Search Options**[View Session History](#)[New Search](#)**Modify Search**

 Check to search only within this results setDisplay Format: Citation Citation & Abstract**» Key**

IEEE JNL IEEE Journal or Magazine

Select Article Information

IEE JNL IEE Journal or Magazine

1. **StateSnap: a snapshot-based interface for state-reproducible operation appliances**
 Kotake, N.; Iwamoto, T.; Suzuki, G.; Aoki, S.; Takashio, K.; Tokuda, H.;
 Mobile and Ubiquitous Systems: Networking and Services, 2005. MobiQuitous
 Second Annual International Conference on
 17-21 July 2005 Page(s):443 - 453
 Digital Object Identifier 10.1109/MOBQUITOUS.2005.55
[AbstractPlus](#) | Full Text: [PDF\(792 KB\)](#) IEEE CNF

IEEE CNF IEEE Conference Proceeding

2. **Electron images of vidicon cathodes**
 Corson, B.R.;
 Electron Devices, IEEE Transactions on
 Volume 12, Issue 8, Aug 1965 Page(s):449 - 457
[AbstractPlus](#) | Full Text: [PDF\(3224 KB\)](#) IEEE JNL

IEE CNF IEE Conference Proceeding

3. **A logic programming framework for modeling temporal objects**
 Kesim, F.N.; Sergot, M.;
 Knowledge and Data Engineering, IEEE Transactions on
 Volume 8, Issue 5, Oct. 1996 Page(s):724 - 741
 Digital Object Identifier 10.1109/69.542026
[AbstractPlus](#) | References | Full Text: [PDF\(1868 KB\)](#) IEEE JNL

IEEE STD IEEE Standard

4. **Low-cost checkpointing and failure recovery in mobile computing system**
 Prakash, R.; Singhal, M.;
 Parallel and Distributed Systems, IEEE Transactions on
 Volume 7, Issue 10, Oct. 1996 Page(s):1035 - 1048
 Digital Object Identifier 10.1109/71.539735
[AbstractPlus](#) | References | Full Text: [PDF\(1580 KB\)](#) IEEE JNL

5. **Independently updated views**
 Kulkarni, U.R.; Ramirez, R.G.;
 Knowledge and Data Engineering, IEEE Transactions on
 Volume 9, Issue 5, Sept.-Oct. 1997 Page(s):798 - 812
 Digital Object Identifier 10.1109/69.634756
[AbstractPlus](#) | References | Full Text: [PDF\(688 KB\)](#) IEEE JNL

6. **On change diagnosis in evolving data streams**

Aggarwal, C.C.;
Knowledge and Data Engineering, IEEE Transactions on
Volume 17, Issue 5, May 2005 Page(s):587 - 600
Digital Object Identifier 10.1109/TKDE.2005.78
[AbstractPlus](#) | Full Text: [PDF\(1400 KB\)](#) IEEE JNL

- 7. A similarity-aware multiversion concurrency control and updating algorithm for snapshots of data**
Gustafsson, T.; Hallqvist, H.; Hansson, J.;
Real-Time Systems, 2005. (ECRTS 2005). Proceedings. 17th Euromicro Conference on Real-Time Systems, 2005. 6-8 July 2005 Page(s):229 - 238
Digital Object Identifier 10.1109/ECRTS.2005.4
[AbstractPlus](#) | Full Text: [PDF\(312 KB\)](#) IEEE CNF
- 8. Byteprints: a tool to gather digital evidence**
Sitaraman, S.; Krishnamurthy, S.; Venkatesan, S.;
Information Technology: Coding and Computing, 2005. ITCC 2005. International Conference on, Volume 1, 4-6 April 2005 Page(s):715 - 720 Vol. 1
Digital Object Identifier 10.1109/ITCC.2005.99
[AbstractPlus](#) | Full Text: [PDF\(200 KB\)](#) IEEE CNF
- 9. Snapshots for semantic maps**
Nielsen, C.W.; Ricks, B.; Goodrich, M.A.; Bruemmer, D.; Few, D.; Few, M.;
Systems, Man and Cybernetics, 2004 IEEE International Conference on, Volume 3, 10-13 Oct. 2004 Page(s):2853 - 2858 vol.3
Digital Object Identifier 10.1109/ICSMC.2004.1400765
[AbstractPlus](#) | Full Text: [PDF\(890 KB\)](#) IEEE CNF
- 10. A Web based history tool for multicast e-meeting sessions**
Parviainen, R.; Parries, P.;
Multimedia and Expo, 2004. ICME '04. 2004 IEEE International Conference on, Volume 1, 27-30 June 2004 Page(s):511 - 514 Vol.1
[AbstractPlus](#) | Full Text: [PDF\(681 KB\)](#) IEEE CNF
- 11. Parallel generation of base relation snapshots for materialized view main warehouse environment**
Saeki, S.; Bhalla, S.; Hasegawa, M.;
Parallel Processing Workshops, 2002. Proceedings. International Conference on, 18-21 Aug. 2002 Page(s):383 - 390
Digital Object Identifier 10.1109/ICPPW.2002.1039755
[AbstractPlus](#) | Full Text: [PDF\(285 KB\)](#) IEEE CNF
- 12. Class-based delta-encoding: a scalable scheme for caching dynamic Web pages**
Psounis, K.;
Distributed Computing Systems Workshops, 2002. Proceedings. 22nd International Conference on, 2-5 July 2002 Page(s):799 - 805
Digital Object Identifier 10.1109/ICDCSW.2002.1030866
[AbstractPlus](#) | Full Text: [PDF\(405 KB\)](#) IEEE CNF
- 13. Optimize CDMA system capacity with location**
Lee, D.J.Y.; Lee, W.C.Y.;
Communications, Computers and signal Processing, 2001. PACRIM. 2001 IEEE International Conference on, Volume 1, 26-28 Aug. 2001 Page(s):17 - 21 vol.1
Digital Object Identifier 10.1109/PACRIM.2001.953512
[AbstractPlus](#) | Full Text: [PDF\(376 KB\)](#) IEEE CNF

- 14. Optimize CDMA system capacity with location**
Lee, W.C.Y.; Lee, D.J.Y.;
Vehicular Technology Conference, 2001. VTC 2001 Fall. IEEE VTS 54th
Volume 2, 7-11 Oct. 2001 Page(s):1015 - 1019 vol.2
Digital Object Identifier 10.1109/VTC.2001.956927
[AbstractPlus](#) | Full Text: [PDF\(320 KB\)](#) IEEE CNF

- 15. A new memory model for selective perception systems**
Soyer, C.; Bozma, H.I.; Istrateanopoulos, Y.;
Intelligent Robots and Systems, 2000. (IROS 2000). Proceedings. 2000 IEEE/
Conference on
Volume 3, 31 Oct.-5 Nov. 2000 Page(s):2304 - 2309 vol.3
Digital Object Identifier 10.1109/IROS.2000.895312
[AbstractPlus](#) | Full Text: [PDF\(500 KB\)](#) IEEE CNF

- 16. Network engineering 2000**
Lee, D.J.Y.; Lee, W.C.Y.;
Microwave and Millimeter Wave Technology, 2000, 2nd International Conference
2000
14-16 Sept. 2000 Page(s):428 - 431
Digital Object Identifier 10.1109/ICMWT.2000.895712
[AbstractPlus](#) | Full Text: [PDF\(328 KB\)](#) IEEE CNF

- 17. Integrating a digital camera in the home environment: architecture and p**
Bennani, N.;
Multimedia Software Engineering, 2000. Proceedings. International Symposium
11-13 Dec. 2000 Page(s):67 - 70
Digital Object Identifier 10.1109/MMSE.2000.897193
[AbstractPlus](#) | Full Text: [PDF\(424 KB\)](#) IEEE CNF

- 18. An SQL3 snapshot**
Melton, J.;
Data Engineering, 1996. Proceedings of the Twelfth International Conference
26 Feb.-1 March 1996 Page(s):666 - 672
Digital Object Identifier 10.1109/ICDE.1996.492217
[AbstractPlus](#) | Full Text: [PDF\(620 KB\)](#) IEEE CNF

- 19. Implementation of sub-aperture sampling into adaptive beamforming using gradient algorithm**
Mandyam, G.; Ahmed, N.; Srinath, M.D.;
Signals, Systems and Computers, 1995. 1995 Conference Record of the Twer
Conference on
Volume 1, 30 Oct.-2 Nov. 1995 Page(s):731 - 734 vol.1
Digital Object Identifier 10.1109/ACSSC.1995.540646
[AbstractPlus](#) | Full Text: [PDF\(248 KB\)](#) IEEE CNF

- 20. A conceptual view of temporal databases**
Coburn, E.J.;
Applied Computing, 1990., Proceedings of the 1990 Symposium on
5-6 April 1990 Page(s):170 - 173
Digital Object Identifier 10.1109/SOAC.1990.82162
[AbstractPlus](#) | Full Text: [PDF\(292 KB\)](#) IEEE CNF

- 21. Oscillatron-1: a connectionist unified theory of cognition**
Atkins, M.A.;
Neural Networks, 1991. 1991 IEEE International Joint Conference on
18-21 Nov. 1991 Page(s):60 - 65 vol.1
Digital Object Identifier 10.1109/IJCNN.1991.170382

[AbstractPlus](#) | Full Text: [PDF\(364 KB\)](#) IEEE CNF

- 22. An intelligent page store for concurrent transaction and query processing**
Dias, D.M.; Goyal, A.; Parr, F.N.;
Research Issues on Data Engineering, 1992: Transaction and Query Processing International Workshop on
2-3 Feb. 1992 Page(s):12 - 19
Digital Object Identifier 10.1109/RIDE.1992.227429
[AbstractPlus](#) | Full Text: [PDF\(748 KB\)](#) IEEE CNF

- 23. Neural network recognition of human face images stored in the database**
Badal, D.Z.;
Computers and Communications, 1993., Twelfth Annual International Phoenix
23-26 March 1993 Page(s):552 - 558
Digital Object Identifier 10.1109/PCCC.1993.344529
[AbstractPlus](#) | Full Text: [PDF\(696 KB\)](#) IEEE CNF

- 24. An efficient subspace algorithm for 2-D harmonic retrieval**
Vanpoucke, F.; Moonen, M.; Berthoumieu, Y.;
Acoustics, Speech, and Signal Processing, 1994. ICASSP-94., 1994 IEEE Int'l Conference on
Volume iv, 19-22 April 1994 Page(s):IV/461 - IV/464 vol.4
Digital Object Identifier 10.1109/ICASSP.1994.389780
[AbstractPlus](#) | Full Text: [PDF\(252 KB\)](#) IEEE CNF



[Help](#) [Contact Us](#) [Privacy & Terms](#)
© Copyright 2005 IEEE -

Indexed by
 Inspec

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 [Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)Results for "((metadata<in>metadata) <and> (inode<in>metadata))" [e-mail](#)

Your search matched 1 of 1286976 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance in Descending** order.» [Search Options](#)[View Session History](#)[Modify Search](#)[New Search](#) » Check to search only within this results setDisplay Format: Citation Citation & Abstract» [Key](#)

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

1. **MRAMFS: a compressing file system for non-volatile RAM**
Edel, N.K.; Tuteja, D.; Miller, E.L.; Brandt, S.A.;
Modeling, Analysis, and Simulation of Computer and Telecommunications Sys
(MASCOTS 2004). Proceedings. The IEEE Computer Society's 12th Annual In
Symposium on
4-8 Oct. 2004 Page(s):596 - 603
Digital Object Identifier 10.1109/MASCOT.2004.1348317
[AbstractPlus](#) | Full Text: PDF(307 KB) IEEE CNF

SEARCH RESULTS[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2005 IEEE -

Indexed by
 Inspec®